Subject: Re: Reading netcdf/hdf5 files slows down
Posted by khyde on Fri, 02 Oct 2015 12:49:47 GMT
View Forum Message <> Reply to Message

On Thursday, October 1, 2015 at 4:45:17 PM UTC-4, Paul van Delst wrote:
> Hello,
>
> On 10/01/15 15:12, KH wrote:
>> Hello,
>>
>> I have several thousand netcdf and hdf5 files that I need to open
>> andextract data from. As the program loops through the files, it
>> progressively takes longer and longer to read each file. Over the course
>> of 24 hours, it went from taking ~6 seconds to completely read the file
>> to more than a minute. Since seconds add up when reading 50,000+ files,
>> I was hoping someone might know why the program would be slowing down
>> and if there is anything (other than restarting IDL) I can do to fix it.
>>
>> For the netcdf files I'm using the following procedures:
>> NCDF_OPEN
>> NCDF_INQUIRE
>> NCDF_ATTNAME
>> NCDF_ATTGET
>> NCDF_CLOSE
>>
>> And for the HDF5:
>> H5F_OPEN
>> H5G_OPEN
>> H5D_OPEN
>> H5D_READ
>> H5D_CLOSE
>> H5G_CLOSE
>> H5F_CLOSE
>>
>> Lastly, I am being sure to close every file once I am done with it
>> andremoving any unneeded variables from memory.
>
> Pawing through the HDF5 docs I see there is also a H5S_CLOSE procedure:
> http://www.exelisvis.com/docs/h5s_close.html
> where it states:
> "Failure to release a dataspace using this procedure will result in
> resource leaks."
>
> Do you create dataspaces and not release them?
>
> cheers,
>
> paulv

Hello,

For every H5x_OPEN I have a corresponding H5x_CLOSE. In my case I'm using H5F, H5G and H5D. I just added H5_CLOSE to the end and I'm hoping that will fix the problem.

Kim
